arXiv:1410.7004v1 [q-bio.PE] 26 Oct 2014

Exact reconciliation of undated trees

Leo van Iersel [*1], Celine Scornavacca[†2], and Steven Kelk[‡3]

[1]Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
[2]ISEM, CNRS – Université Montpellier II, Montpellier, France
[3]Department of Knowledge Engineering (DKE), Maastricht University, Maastricht, The Netherlands

October 28, 2014

Abstract

Reconciliation methods aim at recovering macro evolutionary events and at localizing them in the species history, by observing discrepancies between gene family trees and species trees. In this article we introduce an Integer Linear Programming (ILP) approach for the NP-hard problem of computing a most parsimonious time-consistent reconciliation of a gene tree with a species tree when dating information on speciations is not available. The ILP formulation, which builds upon the $\mathbb{DTL}$ model, returns a most parsimonious reconciliation ranging over all possible datings of the nodes of the species tree. By studying its performance on plausible simulated data we conclude that the ILP approach is significantly faster than a brute force search through the space of all possible species tree datings. Although the ILP formulation is currently limited to small trees, we believe that it is an important proof-of-concept which opens the door to the possibility of developing an exact, parsimony based approach to dating species trees. The software (ILPEACE) is freely available from: `http://homepages.cwi.nl/~iersel/ilpeace/`

Keywords: reconciliation, integer linear programming, dating phylogenies.

1 Background

Reconciliation methods aim at recovering macro evolutionary events – such as speciations, losses and duplications of genes – and at locating them in the species history, by comparing gene family trees to species trees (for a review, see [1]). These methods are often used when studying genome evolution as well as for inferring orthology relationships [2, 3] and for improving phylogeny inference and dating [4, 5, 6].

Here, we consider the problem of finding the Most Parsimonious Reconciliation (MPR) when considering – as possible macro-events that shape the genome – speciations, duplications, transfers and losses of genes. The general problem of finding a MPR is known to

[*]l.j.j.v.iersel@gmail.com. Leo van Iersel was supported by a Veni grant of The Netherlands Organisation for Scientific Research (NWO)

[†]celine.scornavacca@univ-montp2.fr

[‡]steven.kelk@maastrichtuniversity.nl

be NP-complete, even for reconciling two binary trees [7]. The complexity of the problem is due to the difficulty of ensuring the *time-consistency* of gene transfers, i.e. handling the chronological constraints among nodes of the species tree that are induced by transfer events. However, the problem becomes polynomially solvable when accepting a dated species tree as input [8, 9, 10, among others]. In this paper, we extend the combinatorial reconciliation model introduced by Doyon *et al.* [9] (called the "$\mathbb{DTL}$" model), which can be used to solve this special case of the problem (i.e. when the species tree is dated). Although relative dates – obtainable by relaxed molecular clock techniques – are sufficient to ensure tractability, this information is not available for all portions of the Tree of Life. The question, therefore, is how to deal with the NP-hardness of the general (i.e. undated) version of the problem.

We note here that solving the undated version of the problem without ensuring the time-consistency of gene transfers can be done in polynomial time [11]. Moreover, algorithms to solve the undated version of the problem are fast and often find temporally feasible solutions [7, 12, among others]. However, what should be done if the optimal solutions returned by such algorithms are *not* temporally feasible and/or if optimal solutions under the temporally feasible model are strictly less parsimonious than their infeasible counterparts? In this case algorithms that do not enforce temporal feasibility convey only limited information (especially when the algorithm does not indicate if a produced solution is temporally feasible or not). An alternative algorithmic approach is to heuristically search through the space of temporally feasible solution. This is the approach taken by (amongst others) [8]. However, such heuristics offer no guarantees that they will locate the most parsimonious solution. Without such guarantees the MPR model is weakened dramatically, because no indication is given on how far we are from the most parsimonious solution. Ideally, therefore, we require an algorithm that is *guaranteed* to compute the most parsimonious temporally feasible solution.

In this paper, we propose such an algorithm. In particular, we present a flexible Integer Linear Programming (ILP) formulation for finding an MPR when some or all of the dates are unknown. Essentially, the ILP formulation computes an MPR ranging over all possible datings of the species tree. Given that the ILP is built upon the model of Doyon *et al.*, the time-consistency of the reconciliation computed by the ILP is guaranteed. Although ILP has been used earlier in the reconciliation literature [13, 14], this is the first attempt to tackle the $\mathbb{DTL}$ reconciliation model using this technique.

We have embedded the ILP formulation in the software package ILPEACE and made this publicly available [15]. The software computes the MPR between a given binary gene tree and a given undated (or partially dated) binary species tree and outputs a visualisation of the optimal reconciliation in the form of a phylogenetic network and a description of all reconciliation events in a format compatible with the reconciliation editor SylvX [16]. To validate the method (both in terms of correctness and running time) we have compared its performance, on plausible simulated data, to an algorithm that simply brute forces over all possible datings of the species tree. Such a comparison is reasonable because with a mathematical model as complex as the $\mathbb{DTL}$ model it is far from obvious how one can move beyond brute force i.e. how one can intelligently prune the search space. Indeed, this is a strong motivation for our use of ILP in the first place. Our experiments show that ILPEACE is typically 10-100 times faster than the brute force approach and, although it is still limited to relatively small trees, we believe that ILPEACE is nevertheless an important proof-of-concept, mirroring the emergence of proof-of-concept ILP formulations elsewhere in phylogenetics, e.g. [17]. Enhancements to the ILP formulation are likely to open the door to the exciting possibility of, in the future, using a set of gene trees to impose a "most parsimonious dating"

upon an undated species tree, similar to recently undertaken work in the maximum likelihood framework [18].

2 Methods

2.1 Basics

The node set, edge set, internal node set and leaf node set of a tree T are respectively denoted $V(T)$, $E(T)$, $I(T)$ and $L(T)$. Moreover, the label of each leaf u is denoted by $\mathcal{L}(u)$, while the set of labels of leaves of T is denoted by $\mathcal{L}(T)$. The root node of T is denoted by $\rho(T)$. Given two nodes u and v of a rooted tree T, we write $u \leq_T v$ if and only if v is on the unique path from u to the root of T. If $u \leq_T v$ and $u \neq v$ then we write $u <_T v$. For an internal node u of T with two children, let u_l, u_r denote the two children (in arbitrary order). In this paper, we assume that gene and species trees are rooted, binary and uniquely leaf-labeled i.e. within each tree there is a bijection between leaves and labels. Due to this bijectivity we will often refer to leaves and labels interchangeably. The height of a node u in a tree T is denoted by $h_T(u)$ while the height of T is denoted by $h(T)$.

We define a gene tree G as a tree where each leaf represents an extant gene. Similarly, a species tree S is defined as a tree in which each leaf represents a distinct extant species. Note that several leaves of a gene tree can be associated to the same species due to duplication and transfer events. Formally, we indicate this by a surjective function $s : \mathcal{L}(G) \to \mathcal{L}(S)$, called the *species labeling* of G, see Fig. 1 for an example. The set of species labels of the leaves of G is denoted $\mathcal{S}(G)$.

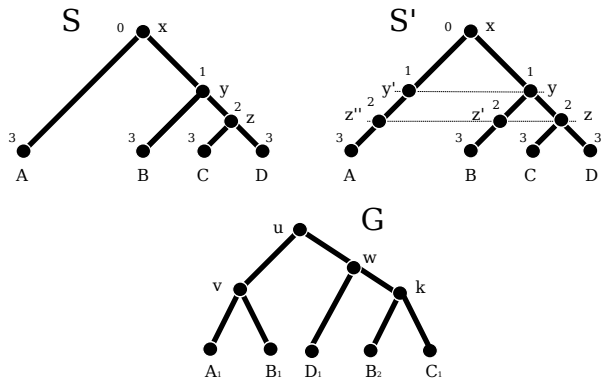

Figure 1: An example of a dated species tree (S) along with its subdivision (S'), and of a gene tree (G). The numbering of nodes of S and S' corresponds to dates. The species labeling of G is the following: $s(A_1) = A$, $s(B_1) = s(B_2) = B$, $s(C_1) = C$ and $s(D_1) = D$.

A species tree S is said to be *dated* if it is associated with a function $\theta_S : V(S) \to \mathbb{R}^+$ such that $\theta_S(\rho(S)) = 0$ and if $y \geq_S x$ then $\theta_S(y) \geq \theta_S(x)$[1]. As mentioned in the introduction, dates make it possible to solve the MPR problem in polynomial time. The recently proposed $\mathbb{DTL}$ model [9] uses a *subdivision* of S rather than S itself when computing reconciliations to ensure time-consistency of gene transfers in polynomial time. Given a tree S and a time function

[1]Note that in the original definition of dating [9], dates are assumed to decrease towards the leaves. This is only a superficial difference.

θ_S, the subdivision S' of S (together with a new time function $\theta_{S'}$) is constructed as follows: firstly, for each node $x \in I(S)$ and each edge $(y_p, y) \in E(S)$ s.t. $\theta_S(y_p) < \theta_S(x) < \theta_S(y)$, an *artificial* node w (i.e. a node with only one child) is inserted along the edge (y_p, y), with $\theta_{S'}(w) = \theta_S(x)$; secondly, for nodes $x \in V(S')$ corresponding to nodes already present in S, we set $\theta_{S'}(x) = \theta_S(x)$. An example of a dated species tree along with its subdivision is shown in Fig. 1.

The $\mathbb{DTL}$ model [9] reconciles a dated binary species tree S with a binary gene tree G by building a mapping α that maps each node $u \in V(G)$ to an ordered list of nodes in $V(S')$, namely the ancestral and/or extant species in which the sequence corresponding to u evolved. This model takes into account four kinds of biological events: speciations, duplications, transfers and losses of genes. The atomic events of this model are: a speciation ($\mathbb{S}$), a duplication ($\mathbb{D}$), a transfer ($\mathbb{T}$), a transfer followed immediately by the loss of the non-transferred child ($\mathbb{TL}$), a speciation followed by the loss of one of the two resulting children ($\mathbb{SL}$), and a contemporary event ($\mathbb{C}$) that associates an extant gene to its corresponding species. Finally, a no event ($\varnothing$), is used to model the fact that a gene lineage has crossed a time boundary. Note that duplication-loss events, unlike transfer-loss and speciation-loss events, leave no trace and are therefore undetectable. This is why, in the $\mathbb{DTL}$ model, losses are never considered alone. The formal definition of a $\mathbb{DTL}$ reconciliation [9] is given below:

Definition 1 ([9]). *Consider a gene tree G, a dated species tree S such that $\mathcal{S}(G) \subseteq \mathcal{L}(S)$, and its subdivision S'. Let α be a function that maps each node u of G onto an ordered sequence of nodes of S', denoted $\alpha(u) = (\alpha_1(u), \alpha_2(u), \ldots, \alpha_\ell(u))$. The function α is said to be a* reconciliation *between G and S' if and only if exactly one of the following events occurs for each couple of nodes u of G and $\alpha_i(u)$ of S' (denoting $\alpha_i(u)$ by x' below):*

a) if x' is the last node of $\alpha(u)$, one of the cases below is true:

1. *$u \in \mathcal{L}(G)$, $x' \in L(S')$ and $s(\mathcal{L}(u)) = \mathcal{L}(x')$;* ($\mathbb{C}$ *event*)
2. *$\{\alpha_1(u_l), \alpha_1(u_r)\} = \{x'_l, x'_r\}$;* ($\mathbb{S}$ *event*)
3. *$\alpha_1(u_l) = x'$ and $\alpha_1(u_r) = x'$;* ($\mathbb{D}$ *event*)
4. *$\alpha_1(u_l) = x'$, and $\alpha_1(u_r)$ is any node other than x' having date $\theta_{S'}(x)$ or $\alpha_1(u_r) = x'$, and $\alpha_1(u_l)$ is any node other than x' having date $\theta_{S'}(x')$;* ($\mathbb{T}$ *event*)

b) otherwise, one of the cases below is true:

5. *x' is an artificial node and $\alpha_{i+1}(u)$ is its only child;* ($\varnothing$ *event*)
6. *x' is not artificial and $\alpha_{i+1}(u) \in \{x'_l, x'_r\}$;* ($\mathbb{SL}$ *event*)
7. *$\alpha_{i+1}(u)$ is any node other than x' having date $\theta_{S'}(x)$.* ($\mathbb{TL}$ *event*)

Note that, even if $\mathbb{D}$, $\mathbb{T}$ and $\mathbb{TL}$ events occur along a branch of S', the mapping is done on nodes of S'. The cost of a reconciliation $c(\alpha)$ is the sum $c_{\mathbb{D}}d + c_{\mathbb{T}}t + c_{\mathbb{L}}\mathrm{l}$, where d is the number of $\mathbb{D}$ events, t is the number of $\mathbb{T}$ and $\mathbb{TL}$ events, and l is the number of $\mathbb{SL}$ and $\mathbb{TL}$ events in α, and $c_{\mathbb{D}}$, $c_{\mathbb{T}}$, $c_{\mathbb{L}}$ are, respectively, the costs of duplications, transfers and losses.

Example 1 An example of reconciliation is depicted in Fig. 2. This reconciliation corresponds to the following mapping α: $\alpha(u) = \{x\}$ (event $\mathbb{S}_1$), $\alpha(v) =$

$\{y'\}$ (event $\mathbb{D}_2$), $\alpha(A_1) = \{y', z'', A\}$ (events $\varnothing_3$, $\varnothing_7$ plus a $\mathbb{C}$ event), $\alpha(B_1) = \{y', z'', z', B\}$ (events $\varnothing_4$, $\mathbb{TL}_6$ $\varnothing_8$ plus a $\mathbb{C}$ event), $\alpha(w) = \{y, z\}$ (events $\mathbb{SL}_5$, $\mathbb{S}_9$), $\alpha(k) = \{C\}$ (event $\mathbb{T}_{10}$), $\alpha(D_1) = \{D\}$ (a $\mathbb{C}$ event), $\alpha(B_2) = \{B\}$ (a $\mathbb{C}$ event), and $\alpha(C_1) = \{C\}$ (a $\mathbb{C}$ event).

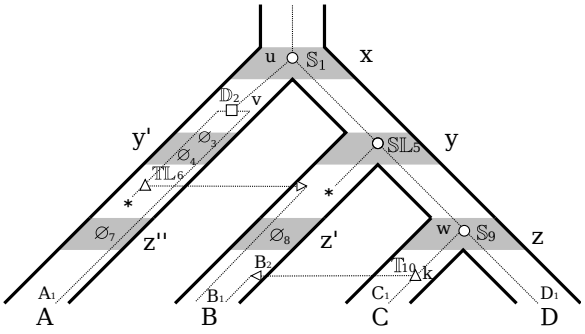

Figure 2: An example of reconciliation for the trees depicted in Fig. 1 containing two $\mathbb{S}$ events, one $\mathbb{SL}$ event, one $\mathbb{D}$ event, a $\mathbb{T}$ event, a $\mathbb{TL}$ event, four $\varnothing$ events and five $\mathbb{C}$ events ($\mathbb{C}$ events are not indicated) .

2.2 Dated Reconciliations

Definition 1 assumes that the dates of all nodes are known. In particular, the subdivision S' of the species tree S in this definition is based on the known dates of the nodes. To deal with situations where the dates of the species tree are not known, we propose the following modified definition. We use $p(.)$ to refer to the parent of a node.

Definition 2. *Consider a gene tree G and a species tree S such that $\mathcal{S}(G) \subseteq \mathcal{L}(S)$. Let α be a function that maps each node u of G onto an ordered sequence of nodes of S, denoted $\alpha(u) = (\alpha_1(u), \alpha_2(u), \ldots, \alpha_\ell(u))$. Moreover, let τ be a function that assigns an ordered sequence of dates $\tau(u) = (\tau_1(u), \ldots, \tau_\ell(u)) \in \mathbb{R}^\ell$ to each node u of G, and let t be a function that assigns a date $t(s) \in \mathbb{R}^+$ to each node s of S. The triple (α, τ, t) is said to be a* dated reconciliation *for (G, S) if and only if exactly one of the following events occurs for each couple of nodes u of G and $\alpha_i(u)$ of S (denoting $\alpha_i(u)$ by x below):*

a) if x is the last node of $\alpha(u)$, one of the cases below is true:

1. *$u \in L(G)$, $x \in L(S)$, $s(\mathcal{L}(u)) = \mathcal{L}(x)$ and $\tau_\ell(u) = t(x)$;* ($\mathbb{C}$ *event*)
2. *$\{\alpha_1(u_l), \alpha_1(u_r)\} = \{x_l, x_r\}$ and $\tau_\ell(u) = t(x)$;* ($\mathbb{S}$ *event*)
3. *$\alpha_1(u_l) = x$, $\alpha_1(u_r) = x$ and $t(p(x)) < \tau_\ell(u) < t(x)$;* ($\mathbb{D}$ *event*)
4. *$\alpha_1(u_l) = x$, and $y := \alpha_1(u_r)$ is any node other than x such that $t(p(x)) < \tau_\ell(u) < t(x)$ and $t(p(y)) < \tau_\ell(u) < t(y)$, or the same with l and r interchanged;* ($\mathbb{T}$ *event*)

b) otherwise, one of the cases below is true:

5. *$\alpha_{i+1}(u) \in \{x_l, x_r\}$ and $\tau_i(u) = t(x)$;* ($\mathbb{SL}$ *event*)
6. *$\alpha_{i+1}(u)$ is any node other than x such that $t(p(x)) < \tau_i(u) < t(x)$ and*

$t(p(\alpha_{i+1}(u)))<\tau_i(u)<t(\alpha_{i+1}(u))$. ($\mathbb{TL}$ *event*)

In addition,

7. $\tau_\ell(p(u)) < \tau_1(u)$ *and, for each* $i \in \{1, \ldots, |\alpha(u)| - 1\}$, $\tau_i(u) < \tau_{i+1}(u)$;
8. $t(p(s)) < t(s)$, *for each node* s *of* S;
9. $t(\rho(S)) = 0$ *and* $t(s) = |I(S)|$ *for each* $s \in L(S)$.

When the species tree is not required to be ultrametric (which we do require in this paper), restriction 9 can be omitted from the above definition.

The cost of a dated reconciliation $c(\alpha, \tau, t)$ is defined similarly to what is done for a reconciliation.

The main differences between Definition 1 and 2 are that in the latter definition we work with the original species tree S and not with its subdivision, and that a date $t(s)$ is assigned to each node s of S and a date $\tau_i(u)$ to each event. For $\mathbb{C}, \mathbb{S}$ and $\mathbb{SL}$ events, the date of the event is equal to the date of the corresponding node of S. For $\mathbb{D}$ events, the date of the event has to be between the date of the corresponding node of S and the date of its parent. For $\mathbb{T}$ and $\mathbb{TL}$ events, the date of the event has to be in-between the date of the donor node of the transfer and the date of its parent, and the date of the event has to be in-between the date of the recipient node of the transfer and the date of its parent. This, along with constraint 7, forces the dated reconciliation to be time-consistent.

Since this definition does not use a subdivision of the species tree but works with the original species tree, it does not use $\varnothing$-events. The core of the idea is that the reconciliation itself indicates the dates of nodes of the species tree and also the dates on which events happen.

The following lemmas ensure that dated reconciliations can be used to find a minimum cost reconciliation over all possible datings of the species tree.

Lemma 1. *Given a reconciliation* α *for a gene tree* G *and a dated species tree* (S, θ_S), *there exists a dated reconciliation* (α', τ, t) *for* G *and* S *with cost* $c(\alpha)$.

Sketch of the proof. Let $(S', \theta_{S'})$ be the dated subdivision of the dated species tree (S, θ_S). The mapping $\alpha'(u)$ for each node $u \in V(G)$ can be obtained from $\alpha(u)$ by removing all the $\varnothing$ events; then the dates of the events can be obtained by setting $\tau_i(u)$ to $\theta_{S'}(\alpha'_i(u))$ for each node $u \in V(G)$ and $1 \leq i \leq |\alpha'(u)|$. To obtain t, we set $t(x) = \theta_S(x)$ for each node $x \in V(S)$. It is straightforward to see that (α', τ, t) satisfies Definition 2 and that it has the same number of events of each type as α. □

Lemma 2. *Given a dated reconciliation* (α', τ, t) *for a gene tree* G *and a species tree* S, *there exists a reconciliation* α *for* G *and* S, t *with cost* $c(\alpha', \tau, t)$.

Sketch of the proof. Constraints 8 and 9 of Definition 2 ensure that the pair (S, t) is a dated tree. Let $(S', \theta_{S'})$ be the dated subdivision of the dated species tree. The mapping $\alpha(u)$ for each node $u \in V(G)$ can be obtained from $\alpha'(u)$ in the following way: we add before $\alpha'_1(u)$ all the descendants x' of $p(\alpha'_1(u))$ in S' having $t_{|\alpha'(u_p)|}(u_p) < \theta_{S'}(x') < t_1(u)$. Moreover, for each $1 \leq i < |\alpha'_i(u)|$, we add between $\alpha'_i(u)$ and $\alpha'_{i+1}(u)$ all the descendants x' of $p(\alpha'_i(u))$ in S' having $t_i(u) < \theta_{S'}(x') < t_{i+1}(u)$.

It can be easily proved that α satisfies Definition 1 and that it has the same number of events of each type as (α', τ, t). □

Example 2

Consider the reconciliation from Fig. 2. Since we do not use a subdivision, the reconciliation is given by the following mapping:

$\alpha(u) = \{x\}$ ($\mathbb{S}$ event) $\quad \alpha(v) = \{A\}$ ($\mathbb{D}$ event)
$\alpha(A_1) = \{A\}$ ($\mathbb{C}$ event) $\quad \alpha(B_1) = \{A, B\}$ ($\mathbb{TL}$ and $\mathbb{C}$ event)
$\alpha(w) = \{y, z\}$ ($\mathbb{SL}$ and $\mathbb{S}$ event) $\quad \alpha(k) = \{C\}$ ($\mathbb{T}$ event)
$\alpha(D_1) = \{D\}$, $\alpha(B_2) = \{B\}$, and $\alpha(C_1) = \{C\}$ (all $\mathbb{C}$ events)

One possible assignment for τ and t is:

$t(x) = 0 \quad t(y) = 1 \quad t(z) = 2 \quad t(A) = (B) = t(C) = (D) = 3$
$\tau(u) = \{0\} \quad \tau(v) = \{0.5\} \quad \tau(A_1) = \{3\} \quad \tau(B_1) = \{1.5, 3\}$
$\tau(w) = \{1, 2\} \quad \tau(k) = \{2.5\} \quad \tau(D_1) = \tau(B_2) = \tau(C_1) = \{3\}.$

Based on this new definition we propose an ILP formulation for finding an optimal dated reconciliation. An ILP consists of three main building blocks: the variables and parameters (described in Section 2.3), the objective function and the constraints (both described in Section 2.4). Although we omit the proof, the ILP formulation is mathematically equivalent to computing a most parsimonious dated reconciliation. (More details on why the two formulations are equivalent will be given in the next section). We call this new method ILPEACE.

2.3 Variables and parameters

Let $\mathcal{I} := \{1, \ldots, \ell\}$, with $\ell = |I(S)|$; note that this is a safe upper bound on the number of nodes of the species tree to which a node of the gene tree can be mapped in a most parsimonious reconciliation (because otherwise one could modify the reconciliation to obtain one with fewer $\mathbb{SL}$ and/or fewer $\mathbb{TL}$ events). Let $\mathbb{E} := \{\mathbb{C}, \mathbb{S}, \mathbb{D}, \mathbb{T}, \mathbb{SL}, \mathbb{TL}\}$. Then we have a binary variable $\alpha_{u,x,i,e}$, for each $u \in V(G), x \in V(S), i \in \mathcal{I}, e \in \mathbb{E}$, and $\alpha_{u,x,i,e} = 1$ precisely if node x is the i-th node of $V(S)$ that node u is mapped to and this mapping corresponds to an event of type e. In addition, there is a variable $\tau_{u,i} \in \mathbb{R}^+$, for each $u \in V(G)$ and $i \in \mathcal{I}$, representing the time of the i-th event mapping node u to some node in $V(S)$ (with $\tau_{u,i} = \tau_{u,i-1}$ if there is no such event). There is also a variable $t_s \in \{0, 1, \ldots, T_{\max}\}$, for each $x \in V(S)$, representing the time of (the speciation event indicated by) node x, with $T_{\max} := |I(S)|$. By assumption, $t_\rho(S) = 0$ and $t_x = T_{\max}$ for each leaf x of the species tree.

The only parameters are the costs of the events. For each event $e \in \mathbb{E}$, c_e indicates its associated cost.

Example 3

Consider the dated reconciliation given in Example 2. Writing the mapping described by this reconciliation as binary variables, we get:

$\alpha_{u,x,1,\mathbb{S}} = 1 \quad \alpha_{v,A,1,\mathbb{D}} = 1 \quad \alpha_{A_1,A,1,\mathbb{C}} = 1$
$\alpha_{B_1,A,1,\mathbb{TL}} = 1 \quad \alpha_{B_1,B,2,\mathbb{C}} = 1$
$\alpha_{w,y,1,\mathbb{SL}} = 1 \quad \alpha_{w,z,2,\mathbb{S}} = 1 \quad \alpha_{k,C,1,\mathbb{T}} = 1$
$\alpha_{D_1,D,1,\mathbb{C}} = 1 \quad \alpha_{B_2,B,1,\mathbb{C}} = 1 \quad \alpha_{C_1,C,1,\mathbb{C}} = 1$

and all other α-variables are 0. The timing variables are:

$t_x = 0$ $t_y = 1$ $t_z = 2$
$t_A = 3$ $t_B = 3$ $t_C = 3$ $t_D = 3$
$\tau_{u,1\leq i\leq 3} = 0$ $\tau_{v,1\leq i\leq 3} = 0.5$
$\tau_{A_1,1\leq i\leq 3} = 3$ $\tau_{B_1,1} = 1.5$ $\tau_{B_1,2\leq i\leq 3} = 3$
$\tau_{w,1} = 1$ $\tau_{w,2\leq i\leq 3} = 2$ $\tau_{k,1\leq i\leq 3} = 2.5$
$\tau_{D_1,1\leq i\leq 3} = 3$ $\tau_{B_2,1\leq i\leq 3} = 3$ $\tau_{C_1,1\leq i\leq 3} = 3.$

We argue that any dated reconciliation can be written in terms of such variables. To see this, let (α, τ, t) be a dated reconciliation. Suppose that $\alpha_i(u) = x$. Then, by Definition 2, there is exactly one event $e \in \mathbb{E}$ corresponding to u, x and i. In that case, we set the binary variable $\alpha_{u,x,i,e}$ to 1 and we set $\tau_{u,i} := \tau_i(u)$. Now suppose that $\alpha_i(u)$ does not exist, i.e. $i > |\alpha(u)|$. In that case, we set $\alpha_{u,x,i,e}$ to 0 and $\tau_{u,i}$ to $\tau_{u,i-1}$. In all other cases, i.e. if $\alpha_i(u) \neq x$, we also set $\alpha_{u,x,i,e}$ to 0. Finally, the dating variables for the species tree are simply obtained by setting $t_x = t(x)$ for each vertex x of S.

The constraints in the next subsection will enforce that the values of the binary variables produced by an ILP solver correspond to a dated reconciliation.

2.4 ILP Formulation

The objective is to minimize the total cost of all events used in the reconciliation. Hence, the objective function is:

$$\min \quad \sum_{u\in V(G)} \sum_{x\in V(S)} \sum_{i\in\mathcal{I}} \sum_{e\in\mathbb{E}} c_e \alpha_{u,x,i,e}.$$

Now we introduce the constraints needed to have our mapping α satisfy Definition 2. Constraints [1]-[7] model some general properties of the mapping. Constraint [1] enforces that each vertex $u \in V(G)$ is mapped to at most one vertex $x \in V(S)$ in at most one type of event $e \in \mathbb{E}$, for each index $i \in \mathcal{I}$. Constraint [2] ensures that each u is mapped somewhere (i.e. to some $x \in V(S)$ for some $e \in \mathbb{E}$) for index $i = 1$. Constraint [3] makes sure that if u is mapped somewhere for index $i + 1$, then u is also mapped somewhere for index i. Constraint [4] states that if u is mapped somewhere for index $i+1$, then for index i it cannot be mapped to an event of type $\mathbb{C}, \mathbb{S}, \mathbb{D}$ or $\mathbb{T}$. Constraint [5] enforces that if u is mapped somewhere for index i with an event of type $\mathbb{SL}$ or $\mathbb{TL}$, then it has to be mapped somewhere for index $i + 1$. Finally, Constraints [6] and [7] state that no events of type $\mathbb{SL}$ and $\mathbb{TL}$ are allowed for the last index $i = \ell$.

[1] $$\sum_{x\in V(S)} \sum_{e\in\mathbb{E}} \alpha_{u,x,i,e} \leq 1 \qquad \forall\ u \in V(G), i \in \mathcal{I}$$

[2] $$\sum_{x\in V(S)} \sum_{e\in\mathbb{E}} \alpha_{u,x,1,e} = 1 \qquad \forall\ u \in V(G)$$

[3] $$\sum_{x\in V(S)} \sum_{e\in\mathbb{E}} (\alpha_{u,x,i,e} - \alpha_{u,x,i+1,e}) \geq 0 \qquad \forall\ u \in V(G), i \in \mathcal{I}\backslash\{\ell\}$$

[4] $$\sum_{x\in V(S)} \sum_{e\in\mathbb{E}} \alpha_{u,x,i+1,e} + \sum_{x\in V(S)} \sum_{e\in\{\mathbb{C},\mathbb{S},\mathbb{D},\mathbb{T}\}} \alpha_{u,x,i,e} \leq 1 \qquad \forall u \in V(G), i \in \mathcal{I}\backslash\{\ell\}$$

$$[5] \quad \sum_{x \in V(S)} \sum_{e \in \{\mathbb{SL}, \mathbb{TL}\}} \alpha_{u,x,i,e} - \sum_{x \in V(S)} \sum_{e \in \mathbb{E}} \alpha_{u,x,i+1,e} \leq 0 \qquad \forall\ u \in V(G), i \in \mathcal{I} \backslash \{\ell\}$$

$$[6] \quad \alpha_{u,x,\ell,\mathbb{SL}} = 0 \qquad \forall\ u \in V(G), x \in V(S)$$

$$[7] \quad \alpha_{u,x,\ell,\mathbb{TL}} = 0 \qquad \forall\ u \in V(G), x \in V(S)$$

The next set of constraints models the $\mathbb{C}$ events. A $\mathbb{C}$ event maps a leaf of G to the leaf of S given by the $s(.)$ function.

$$[8] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{C}} = 1 \qquad \forall\ u \in L(G), x \in L(S) \text{ with } s(\mathcal{L}(u)) = \mathcal{L}(x)$$

$$[9] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{C}} = 0 \qquad \text{otherwise}$$

We now model the $\mathbb{S}$ events. If u is mapped to x in a $\mathbb{S}$ event, then the children of u must be mapped to the children of x. This can be enforced by the following constraints. Constraint [10] states that leaves can not be mapped in $\mathbb{S}$ events. Constraint [11] enforces that if u is mapped somewhere in an event of type $\mathbb{S}$, then its "left" child u_l must be mapped to one of the children x_l, x_r of x for index $i = 1$. Constraint [12] does the same for the right child u_r of u. Constraint [13] then enforces that u_l and u_r cannot both be mapped to x_l and Constraint [14] enforces that u_l and u_r cannot both be mapped to x_r.

$$[10] \quad \alpha_{u,x,i,\mathbb{S}} = 0 \qquad \forall i \in \mathcal{I}, \text{if } u \in L(G) \text{ or } x \in L(S)$$

$$[11] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{S}} - \sum_{e \in \mathbb{E}} (\alpha_{u_l,x_l,1,e} + \alpha_{u_l,x_r,1,e}) \leq 0 \qquad \forall\ u \in I(G), x \in I(S)$$

$$[12] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{S}} - \sum_{e \in \mathbb{E}} (\alpha_{u_r,x_l,1,e} + \alpha_{u_r,x_r,1,e}) \leq 0 \qquad \forall\ u \in I(G), x \in I(S)$$

$$[13] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{S}} + \sum_{e \in \mathbb{E}} (\alpha_{u_l,x_l,1,e} + \alpha_{u_r,x_l,1,e}) \leq 2 \qquad \forall\ u \in I(G), x \in I(S)$$

$$[14] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{S}} + \sum_{e \in \mathbb{E}} (\alpha_{u_l,x_r,1,e} + \alpha_{u_r,x_r,1,e}) \leq 2 \qquad \forall\ u \in I(G), x \in I(S)$$

This brings us to modeling $\mathbb{D}$ events. If u is mapped to x in a $\mathbb{D}$ event, then the children of u must also be mapped to x. This is modelled by the following two constraints. Constraint [15] makes sure that leaves are not mapped in events of type $\mathbb{D}$. Constraint [16] enforces that if u is mapped to x in an event of type $\mathbb{D}$, then both children u_l and u_r of u have to be mapped to x.

$$[15] \quad \alpha_{u,x,i,\mathbb{D}} = 0 \qquad \forall\ u \in L(G), x \in V(S), i \in \mathcal{I}$$

$$[16] \quad \sum_{i \in \mathcal{I}} 2\alpha_{u,x,i,\mathbb{D}} - \sum_{e \in \mathbb{E}} (\alpha_{u_l,x,1,e} + \alpha_{u_r,x,1,e}) \leq 0 \qquad \forall\ u \in I(G), x \in V(S)$$

Constraints [17]-[28] model general properties of the timing variables. First we need some definitions. For a node v that is not the root, let $p(v)$ denote its parent. Let $l_{p(v),v}$ denote

the length of the edge between $p(v)$ and v (representing the elapsed time between the two speciations), if known, and let $l_{p(v),v} = 0$ indicate that the length of this edge is not known. Let $\epsilon := 1/\ell$.

The following five constraints model the timing variables of the vertices $V(S)$ of the species tree. Constraint [17] states that the date of the root is 0. We assume that all leaves of S have date $T_{\max} := |I(S)|$ (Constraint [18]) while the internal vertices of S have dates in $\{0, 1, \ldots, T_{\max} - 1\}$ (Constraint [19]). (Note that these constraints could easily be relaxed for non-ultrametric species trees.) Constraint [20] makes sure that the difference between the date of x and the date of its parent $p(x)$ is $l_{p(x),x}$ if this edge length is known. Constraint [21] states that this difference should be at least 1 when the edge length is not known. This constraint is valid because without loss of generality the dates of internal vertices of S are all different elements of $\{0, 1, \ldots, T_{\max} - 1\}$.

[17] $t_{\rho(S)} = 0$

[18] $t_x = T_{\max}$ $\quad \forall\ x \in L(S)$

[19] $t_x \leq T_{\max} - 1$ $\quad \forall\ x \in I(S)$

[20] $t_x - t_{p(x)} = l_{p(x),x}$ $\quad \forall\ x \in V(S)\backslash\{\rho(S)\}$ with $l_{p(x),x} \neq 0$

[21] $t_x - t_{p(x)} \geq 1$ $\quad \forall\ x \in V(S)\backslash\{\rho(S)\}$ with $l_{p(x),x} = 0$

The next three constraints model the timing variables for the mapping events. Constraint [22] states that the date of the event mapping vertex u for index $i = 1$ should be strictly larger than the date of the last event mapping the parent $p(u)$ of u. Constraints [23] and [24] ensure that $\tau_{u,i+1} \geq \tau_{u,i}$ and that equality holds precisely if u is not mapped anywhere for index $i + 1$.

[22] $\tau_{u,1} - \tau_{p(u),\ell} \geq \epsilon$ $\quad \forall\ u \in V(G)$

[23] $\tau_{u,i+1} - \tau_{u,i} - \sum_{x \in V(S)} \sum_{e \in \mathbb{E}} \alpha_{u,x,i+1,e} \geq 0$ $\quad \forall\ u \in V(G), i \in \mathcal{I}\backslash\{\ell\}$

[24] $\tau_{u,i+1} - \tau_{u,i} - T_{\max} \sum_{x \in V(S)} \sum_{e \in \mathbb{E}} \alpha_{u,x,i+1,e} \leq 0$ $\quad \forall\ u \in V(G), i \in \mathcal{I}\backslash\{\ell\}$

Constraints [25] and [26] below make sure that the time of an event is in-between the time of the corresponding node x of the species tree and the time of its parent. If the event is of type $\mathbb{C}$, $\mathbb{S}$ or $\mathbb{SL}$, then the time of the event must be equal to the time of x (Constraint [27]), and if the event is of type $\mathbb{D}$, $\mathbb{T}$ or $\mathbb{TL}$, then the time of the event must be strictly smaller than the time of x (Constraint [28]), since the latter events happen on edges rather than nodes of the species tree. Note that we do not require here that $\tau_{u,i}$ is strictly greater than $t_{p(x)}$ because this is already implied by Constraint [22]. Let $V^+(S) := V(S)\backslash\{\rho(S)\}$.

[25] $\tau_{u,i} - t_x + T_{\max} \sum_{e \in \mathbb{E}} \alpha_{u,x,i,e} \leq T_{\max}$ $\quad \forall\ u \in V(G), x \in V(S), i \in \mathcal{I}$

[26] $t_{p(x)} - \tau_{u,i} + T_{\max} \sum_{e \in \mathbb{E}} \alpha_{u,x,i,e} \leq T_{\max}$ $\quad \forall\ u \in V(G), x \in V^+(S), i \in \mathcal{I}$

$$[27] \quad t_x - \tau_{u,i} + T_{\max} \sum_{e \in \{\mathbb{C},\mathbb{S},\mathbb{SL}\}} \alpha_{u,x,i,e} \leq T_{\max} \qquad \forall\, u \in V(G), x \in V^+(S), i \in \mathcal{I}$$

$$[28] \quad \tau_{u,i} - t_x + T_{\max} \sum_{e \in \{\mathbb{D},\mathbb{T},\mathbb{TL}\}} \alpha_{u,x,i,e} \leq T_{\max} - \epsilon \qquad \forall\, u \in V(G), x \in V^+(S), i \in \mathcal{I}$$

We can now model $\mathbb{T}$ events. First, Constraint [29] states that leaves cannot be mapped in $\mathbb{T}$ events. Constraint [30] enforces that if u is mapped to x in a $\mathbb{T}$ event, then one child of u must also be mapped to x. Say that the other child of u is mapped to y, then we must have $y \in U(x)$, with $U(x)$ the set of nodes that are not a descendant and not an ancestor of x. This is ensured by Constraints [31] and [32].

$$[29] \quad \alpha_{u,x,i,\mathbb{T}} = 0 \qquad \forall\, u \in L(G), x \in V(S), i \in \mathcal{I}$$

$$[30] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{T}} - \sum_{e \in \mathbb{E}} (\alpha_{u_l,x,1,e} + \alpha_{u_r,x,1,e}) \leq 0 \qquad \forall\, u \in I(G), x \in V(S)$$

$$[31] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{T}} + \sum_{e \in \mathbb{E}} \alpha_{u_l,x,1,e} - \sum_{e \in \mathbb{E}} \sum_{y \in U(x)} \alpha_{u_r,y,1,e} \leq 1 \qquad \forall\, u \in I(G), x \in V(S)$$

$$[32] \quad \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{T}} + \sum_{e \in \mathbb{E}} \alpha_{u_r,x,1,e} - \sum_{e \in \mathbb{E}} \sum_{y \in U(x)} \alpha_{u_l,y,1,e} \leq 1 \qquad \forall\, u \in I(G), x \in V(S)$$

Constraints [33] and [34] enforce that if $c \in \{u_l, u_r\}$ is mapped to y for index $i = 1$, then $t_{p(y)} < \tau_{u,i} < t_y$[2] . If u is mapped by an event of type $\mathbb{T}$ for index $i = \ell$, then this restriction is necessary to enforce time-consistency. Note that, if the last event mapping u is an event not of type $\mathbb{T}$, then a child of u can only be mapped to either y or to a child of y. In both cases, the restriction is valid. Since the restriction is always fulfilled for any event that is not a $\mathbb{T}$, for simplicity we impose it for all event types.

$$[33] \quad \tau_{u,\ell} - t_y + T_{\max} \sum_{e \in \mathbb{E}} \alpha_{c,y,1,e} \leq T_{\max} \qquad \forall\, u \in V(G), c \in \{u_l, u_r\}, y \in V(S)$$

$$[34] \quad t_{p(y)} - \tau_{u,\ell} + T_{\max} \sum_{e \in \mathbb{E}} \alpha_{c,y,1,e} \leq T_{\max} \qquad \forall\, u \in V(G), c \in \{u_l, u_r\}, y \in V^+(S)$$

The next two constraints model $\mathbb{SL}$ events. Constraint [35] states that leaves cannot be mapped in $\mathbb{SL}$ events. Constraint [36] enforces that if u is mapped to x for index i in an $\mathbb{SL}$ event, then u must be mapped to one of the children of x for index $i + 1$.

$$[35] \quad \alpha_{u,x,i,\mathbb{SL}} = 0 \qquad \forall\, u \in V(G), x \in L(S), i \in \mathcal{I}$$

$$[36] \quad \alpha_{u,x,i,\mathbb{SL}} - \sum_{e \in \mathbb{E}} (\alpha_{u,x_l,i+1,e} + \alpha_{u,x_r,i+1,e}) \leq 0 \qquad \forall\, u \in V(G), x \in I(S), i \in \mathcal{I} \backslash \{\ell\}$$

[2] Actually, for keeping the ILP formulation as simple as possible, Constraints [33] and [34] enforce that $t_{p(y)} \leq \tau_{u,\ell} \leq t_y$. If $\tau_{u,\ell}$ is equal to $t_{p(y)}$ (or to t_y), we can simply add (or subtract) $\epsilon/2$ to $\tau_{u,\ell}$ to satisfy the strict inequality.

We now formulate the constraints that model $\mathbb{TL}$ events. If u is mapped to x for index i in a $\mathbb{TL}$-event, and if u is mapped to y for index $i+1$, then we must have that $t_{p(y)} < \tau_{u,i} < t_y$ (see footnote 2). This is enforced by Constraints [37] and [38]. Note that this holds automatically for $\mathbb{SL}$ events, hence we do not need to restrict it to $\mathbb{TL}$ events. Constraint [39] enforces that $y \notin U(x)$.

$$[37] \quad \tau_{u,i} - t_y + T_{\max} \sum_{e \in \mathbb{E}} \alpha_{u,y,i+1,e} \leq T_{\max} \qquad \forall\, u \in V(G), y \in V(S), i \in \mathcal{I} \backslash \{\ell\}$$

$$[38] \quad t_{p(y)} - \tau_{u,i} + T_{\max} \sum_{e \in \mathbb{E}} \alpha_{u,y,i+1,e} \leq T_{\max} \qquad \forall\, u \in V(G), y \in V^+(S), i \in \mathcal{I} \backslash \{\ell\}$$

$$[39] \quad \alpha_{u,x,i,\mathbb{TL}} - \sum_{y \in U(x)} \sum_{e \in \mathbb{E}} \alpha_{u,y,i+1,e} \leq 0 \qquad \forall\, u \in V(G), x \in V(S), i \in \mathcal{I} \backslash \{\ell\}$$

Finally, we need the bounds and integrality constraints. Even though the time-variables could be restricted to be integer, this is not necessary.

$$[40] \quad 0 \leq t_x \leq T_{\max} \qquad \forall\, x \in V(S)$$

$$[41] \quad 0 \leq \tau_{u,i} \leq T_{\max} \qquad \forall\, u \in V(G), i \in \mathcal{I}$$

$$[42] \quad \alpha_{u,x,i,e} \in \{0,1\} \qquad \forall\, u \in V(G), x \in V(S), i \in \mathcal{I}, e \in \mathbb{E}$$

This concludes the ILP formulation.

For example, it can be easily checked that the given assignment of binary variables and dates in Example 3 fulfills all the constraints of the ILP formulation for the trees in Figure 1 and thus is a valid dated reconciliation for these trees.

3 Results and discussion

3.1 Implementation

The ILPEACE method has been implemented in Java and made publicly available [15]. It generates an ILP formulation which is then solved by the ILP solver CPLEX [19], which is a state-of-the-art ILP solver built upon a polyhedral *branch, bound and cut* core [20]. ILPEACE first computes an approximate solution by setting the number of $\mathbb{TL}$ events to 0. This can be achieved by adding a constraint

$$[43] \quad \sum_{u \in V(G)} \sum_{x \in V(S)} \sum_{i \in \mathcal{I}} \alpha_{u,x,i,\mathbb{TL}} = 0.$$

Practical experiments show that an optimal solution to this restricted problem can usually be computed relatively quickly by CPLEX and that it provides a very good upper bound (see Section 3.2.2). We call this the *no-$\mathbb{TL}$-bound*. The corresponding solution is then given to CPLEX as a "warm start" (i.e. an upper bound) from which CPLEX can search for the true optimum. Once an optimal reconciliation has been found, ILPEACE outputs the reconciliation score, number of events of each type, computation time, and the following files: a file containing a subdivision of the species tree and one containing the gene tree (in both files each node is associated with an id), the reconciliation in a format compatible with the

reconciliation editor SylvX [16] and finally a file containing a phylogenetic network in eNewick format [21] obtained by adding transfer edges to the species tree (as indicated by the optimal reconciliation).

Note that, once the optimal dating is return by ILPEACE, a graph containing all optional solutions can be constructed using the software presented in [22], an implementation of which is available at `http://mbb.univ-montp2.fr/MBB/subsection/downloads`. (The ILP solver returns only a single optimal dating, but modern ILP solvers such as CPLEX offer access to "solution pools" from which alternative optima can be sampled. As with many optimization problems there can be exponentially many optima so generation of them all is challenging. However, ILPEACE does of course compute the *score* of the most parsimonious reconcilation, which is an essential first step towards understanding the structure of this optimal space.)

3.2 Validation on simulated data

We validated the ILPEACE method by comparing its performance on simulated data to a brute force approach. Note that, since the aim of this section is to compare the performance of these two methods, we are not concerned here by retrieving the true reconciliation neither the true $\mathbb{DTL}$ events. The comparsion with a brute force approach is reasonable because it is unclear how one could develop, for example, an ad-hoc branch and bound algorithm to prune the search space intelligently; this is because the underlying mathematical model is far more complex than more classical combinatorial optimization problems. (Indeed, this complexity is a major reason why we chose to address the problem with ILP.) We also chose not to compare ILPEACE to the superficially similar RANGER-DTL software from [12] or the JANE algorithm from [8]. RANGER-DTL does not guarantee time-consistency (and does not report whether its optimal solution is time-consistent). JANE uses a slightly different event model to the $\mathbb{DTL}$ model used here and we have observed that JANE sometimes returns strictly less parsiminious solutions than ILPEACE. It is unclear whether this is due to model differences or generation of suboptimal solutions because JANE is a heuristic and offers no formal guarantees that the solutions it finds are optimal (even within its own model).

The brute force approach consists of three steps. Step 1 is to find all possible orderings for the nodes of the undated species tree. This is done by generating all possible linear extensions (i.e. total orders) of the partial ordering implied by the topology of the species tree [23]. Step 2 is to apply, for each thus obtained dated version of the species tree, the $\mathbb{DTL}$ reconciliation method [9, 22]. Step 3 is to return an ordering minimizing the reconciliation cost over all possible orderings of the species tree.

Note that, to try all possible orderings for the species tree nodes, and thus to solve the undated $\mathbb{DTL}$ reconciliation problem, we would in principle have to loop through all orderings that extend the partial ordering given by the species tree, not only the total ones. Luckily, it can be proven that looping through all total orderings is enough to find a most parsimonious dating of S. The full proof has been omitted and we give here a short sketch. For each two nodes s, s' of S such that $\theta_S(s) = \theta_S(s')$, the value of $\theta_S(s')$ can be modified to $\theta_S(s') - \epsilon$ for some small $\epsilon > 0$. If ϵ is chosen small enough – i.e. such that no other date of S falls in the interval $[\theta_S(s') - \epsilon, \theta_S(s')]$ – this can be shown not to negatively affect the reconciliation. Doing so for each pair of nodes with the same value $\theta_S(\cdot)$ leads to a date function θ'_S for which $\theta'_S(s) \neq \theta'_S(s')$ whenever $s \neq s'$. Such a date function can easily be turned into a total ordering.

3.2.1 Gene tree simulation

To simulate gene trees along a dated species tree S, we start by associating to every branch of S an *activity* a that represents the overall rate at which $\mathbb{DTL}$ events occur on this branch, along with specific rates for each individual event type $r_{\mathbb{D}}, r_{\mathbb{T}}, r_{\mathbb{L}}$, with $a = r_{\mathbb{D}} + r_{\mathbb{T}} + r_{\mathbb{L}}$. We then use a birth-and-death process [24] to simulate each gene tree using a scheme similar to what done in [25]:

1. at the beginning of the process, the first gene of the birth-and-death process is located at the root of S;

2. at any time, the time t_{next} of the next potential $\mathbb{DTL}$ event in every existing gene is calculated by simulating an exponential variable with parameter equal to the activity of the branch (x, y) containing that gene. Then, if $t_{next} \geq \theta_{S'}(y)$, the next event is determined to be a $\mathbb{C}$ event if y is a leaf, and an $\mathbb{S}$ event otherwise. If $t_{next} < \theta_{S'}(y)$, the next event is a $\mathbb{DTL}$ event and we rely on the relative rates $r_{\mathbb{D}}, r_{\mathbb{T}}, r_{\mathbb{L}}$ to determine its type.

3. we repeat this process until we reach the time of the extant species.

Only gene trees with at least 2 leaves are retained.

For the simulations, we chose as dated species tree a phylogeny of 37 proteobacteria over a period of 500 million years [26]. To be able to apply the brute force approach, we were forced to run the gene tree simulation using a version of this phylogeny restricted to 17 leaves.

The duplication, transfer and loss rates were generated for each simulated gene independently and were chosen in accordance with real dataset observations [26] using the same scheme as [27]: the loss rate was randomly chosen in the interval [0.001, 0.0018], where the units are events per gene per million years; the ratio between the "birth" rate (sum of the duplication and transfer rates) and the loss rate was randomly chosen in the interval [0.5,1.1]; ultimately the proportion of the duplication rate to the birth rate was randomly chosen in the interval [0.7,1]. We chose to simulate gene trees using a scheme similar to the one used in [25] rather than the one used in [27] to avoid to simulate sequence alignments, computationally costly and not needed in this paper.

We simulated 100 gene trees and we applied both the ILPEACE and the brute force approach to each pair species tree - gene tree, using the following cost vector: $c_{\mathbb{D}} = 2$, $c_{\mathbb{T}} = 3$ and $c_{\mathbb{L}} = 1$ (and $c_{\mathbb{SL}} = 1$, $c_{\mathbb{TL}} = 4$ and $c_{\mathbb{S}} = c_{\mathbb{C}} = 0$). We ran the experiments on a 3.2 GHz Intel Core i3 processor with 8 GB of RAM, using CPLEX 12.5 as the ILP solver.

3.2.2 Results

The results of the simulations are shown in Table 1.

As can be seen from the table, ILPEACE is on average between one and two orders of magnitude faster than the brute force algorithm. Note also that the no-$\mathbb{TL}$-bound, which appears to be significantly easier for CPLEX to compute than the true optimum, is nevertheless an extremely good approximation of the true optimum.

Table 1: Minimum, average and maximum values for the number of taxa in the gene trees, reconciliation scores, running times of the brute force approach, the no-$\mathbb{TL}$-bound and ILPEACE, and the approximation ratio of the no-$\mathbb{TL}$-bound. Running times are in seconds.

			running time			app. ratio no-$\mathbb{TL}$
	taxa	score	brute force	no-$\mathbb{TL}$	ILPEACE	over the true optimum
min.	2	0	384	2.3	6	1
avg.	11.2	12.2	748.1	5.66	18.6	1.0007
max.	20	40	1968	24.3	185.6	1.07

4 Conclusions

We have shown that in the absence of complete dating information for the species tree, ILP can be a powerful tool for computing a most parsimious temporally feasible reconciliation ranging over the space of all possible datings. This is the first algorithm with such properties, since competing approaches either fail to guarantee optimality (i.e. that the solution is most parsimonious) or feasibility (i.e. that the solution is time consistent). For the trees we used (up to 20 taxa) the ILP formulation was typically 10-100 times faster than the obvious brute force algorithm. This is significant given that it is far from obvious how to develop an *ad hoc* pruning algorithm to substantially improve upon the brute force approach. The average running time of the ILP formulation was 18 seconds. For larger trees (say, 30 taxa or more, which is already far out of reach of a naive brute force algorithm) the performance of ILPEACE unfortunately begins to deteriorate. It will still terminate quickly in many cases but worst-case running times start to explode. Nevertheless, we are optimistic that ILPEACE represents an important proof-of-concept, for the following reasons. Firstly, ongoing research into an improved ILP formulation is likely to yield a significant improvement in running times; the most obvious starting point is to understand why $\mathbb{TL}$ events are a bottleneck to fast ILP execution. Indeed, in the combinatorial optimization literature it is standard practice to introduce an initial ILP formulation which then in subsequent publications is steadily refined, see e.g. [28]. Secondly, ILP is a very attractive approach because of its inherent flexibility: extra knowledge (such as partial dating information) can easily be incorporated into the formulation. Moreover, the formulation can easily be extended to allow simultaneous reconciliation of many gene trees with a given species tree. As mentioned in the introduction this will potentially enable us to use several gene trees to impose a "most parsimonious dating" on a single species tree.

Competing interests

The authors declare that they have no competing interests.

Author's contributions

LvI, CS and SK designed the model and solution approach, verified its correctness, designed and conducted the experiments and wrote the paper. LvI implemented the software.

Acknowledgements

The authors would like to thank Eric Tannier and Bastien Boussau for fruitful discussions.

References

[1] Doyon, J.P., Ranwez, V., Daubin, V., Berry, V.: Models, algorithms and programs for phylogeny reconciliation. Briefings in Bioinformatics **12**, 392–400 (2011)

[2] Storm, C.E.V., Sonnhammer, E.L.L.: Automated ortholog inference from phylogenetic trees and calculation of orthology reliability. Bioinformatics **18**(1), 92–99 (2002)

[3] van der Heijden, R., Snel, B., van Noort, V., Huynen, M.: Orthology prediction at scalable resolution by phylogenetic tree analysis. BMC Bioinformatics **8**(1), 83 (2007)

[4] Abby, S., Tannier, E., Gouy, M., Daubin, V.: Detecting lateral gene transfers by statistical reconciliation of phylogenetic forests. BMC Bioinformatics **11**(1), 324 (2010)

[5] Åkerborg, Ö., Sennblad, B., Arvestad, L., Lagergren, J.: Simultaneous Bayesian gene tree reconstruction and reconciliation analysis. Proceedings of the National Academy of Sciences of the United States of America **106**(14), 5714–5719 (2009)

[6] Boussau, B., Szöllősi, G.J., Duret, L., Gouy, M., Tannier, E., Daubin, V.: Genome-scale coestimation of species and gene trees. Genome Research **23**(2), 323–330 (2013)

[7] Tofigh, A., Hallett, M., Lagergren, J.: Simultaneous identification of duplications and lateral gene transfers. IEEE/ACM Transactions on Computational Biology and Bioinformatics **8**(2), 517–535 (2011)

[8] Conow, C., Fielder, D., Ovadia, Y., Libeskind-Hadas, R.: Jane: a new tool for the cophylogeny reconstruction problem. Algorithms for Molecular Biology **5**, 16 (2010)

[9] Doyon, J.-P., Scornavacca, C., Gorbunov, K.Y., Szöllősi, G.J., Ranwez, V., Berry, V.: An efficient algorithm for gene/species trees parsimonious reconciliation with losses, duplications and transfers. In: Proceedings of the 2010 International Conference on Comparative Genomics. RECOMB-CG'10, pp. 93–108. Springer, Berlin, Heidelberg (2011)

[10] Tofigh, A.: Using trees to capture reticulate evolution, lateral gene transfers and cancer progression. PhD thesis, KTH Royal Institute of Technology, Sweden (2009)

[11] Hallett, M., Largergren, J., Tofigh, A.: Simultaneous identification of duplications and lateral transfers. In: Proceedings of the Eight International Conference on Research in Computational Molecular Biology (RECOMB), pp. 347–356 (2004)

[12] Bansal, M.S., Alm, E.J., Kellis, M.: Efficient algorithms for the reconciliation problem with gene duplication, horizontal transfer and loss. Bioinformatics **28**(12), 283–291 (2012)

[13] Chang, W.-C., Burleigh, G.J., Fernández-Baca, D.F., Eulenstein, O.: An ILP solution for the gene duplication problem. BMC Bioinformatics **12**(Suppl 1), 14 (2011)

[14] Than, C., Nakhleh, L.: Species tree inference by minimizing deep coalescences. PLoS Computational Biology **5**(9), 1000501 (2009)

[15] van Iersel, L.J.J., Scornavacca, C., Kelk, S.M.: ILPEACE. `http://homepages.cwi.nl/~iersel/ilpeace/`

[16] Chevenet, F.: SylvX. work in progress, `http://www.sylvx.org`

[17] Dong, J., Fernandez-Baca, D., McMorris, F.: Constructing majority-rule supertrees. Algorithms for Molecular Biology **5**(1), 2 (2010)

[18] Szöllősi, G.J., Boussau, B., Abby, S.S., Tannier, E., Daubin, V.: Phylogenetic modeling of lateral gene transfer reconstructs the pattern and relative timing of speciations. Proceedings of the National Academy of Sciences **109**(43), 17513–17518 (2012)

[19] IBM-ILOG: CPLEX version 12.5 (2013)

[20] Mitchell, J.E.: Branch-and-cut algorithms for combinatorial optimization problems. Handbook of Applied Optimization, 65–77 (2002)

[21] Cardona, G., Rosselló, F., Valiente, G.: Extended Newick: it is time for a standard representation of phylogenetic networks. BMC Bioinformatics **9**, 532 (2008)

[22] Scornavacca, C., Paprotny, W., Berry, V., Ranwez, V.: Representing a set of reconciliations in a compact way. Journal of Bioinformatics and Computational Biology **11** (2013)

[23] Pruesse, G., Ruskey, F.: Generating linear extensions fast. SIAM Journal on Computing **23**(2), 373–386 (1994)

[24] Kendall, D.G.: On the generalized birth-and-death process. Ann Math Stat **19**, 1–15 (1948)

[25] Chan, Y.-b., Ranwez, V., Scornavacca, C.: Reconciliation-based detection of co-evolving gene families. BMC Bioinformatics **14**(332) (2013)

[26] David, L.A., Alm, E.J.: Rapid evolutionary innovation during an archaean genetic expansion. Nature **469**(7328), 93–96 (2011)

[27] Nguyen, T.H., Ranwez, V., Pointet, S., Chifolleau, A.-M.A., Doyon, J.-P., Berry, V.: Reconciliation and local gene tree rearrangement can be of mutual profit. Algorithms for Molecular Biology **8**(1), 12 (2013)

[28] Lancia, G., Serafini, P.: A set-covering approach with column generation for parsimony haplotyping. INFORMS Journal on Computing **21**(1), 151–166 (2009)